

BAYESIAN NOISE ESTIMATION FOR NON-IDEAL CMB EXPERIMENTS

I. K. WEHUS^{1,2}, S. K. NÆSS³ AND H. K. ERIKSEN^{3,4}

Draft version October 7, 2011

ABSTRACT

We describe a Bayesian framework for estimating the time-domain noise covariance of CMB observations, typically parametrized in terms of a $1/f$ frequency profile. This framework is based on the Gibbs sampling algorithm, which allows for exact marginalization over nuisance parameters through conditional probability distributions. In this paper we implement support for gaps in the data streams and marginalization over fixed time-domain templates, and also outline how to marginalize over confusion from CMB fluctuations, which may be important for high signal-to-noise experiments. As a by-product of the method, we obtain proper constrained realizations, which themselves can be useful for map making. To validate the algorithm, we demonstrate that the reconstructed noise parameters and corresponding uncertainties are unbiased using simulated data. The CPU time required to process a single data stream of 100 000 samples with 1000 samples removed by gaps is 3 seconds if only the maximum posterior parameters are required, and 21 seconds if one also want to obtain the corresponding uncertainties by Gibbs sampling.

Subject headings: cosmic microwave background — cosmology: observations — methods: statistical

1. INTRODUCTION

Detailed observations of the cosmic microwave background (CMB) during the last two decades have revolutionized cosmology. Through detailed measurements of the angular CMB power spectrum, a highly successful cosmological concordance model has been established, stating that the universe is statistically isotropic and homogeneous, filled with Gaussian random fluctuations drawn from a Λ CDM spectrum, and consists of 4% baryonic matter, 23% dark matter and 73% dark energy (e.g., Komatsu et al. 2011, and references therein). Using this model, millions of data points from many different types of cosmological observations can be fitted with only six free parameters.

This success has been driven primarily by rapid progress in CMB detector technology, allowing experimentalists to make more and more detailed maps of the CMB fluctuations. However, such maps are imperfect, in the sense that they typically are contaminated by various instrumental effects. For instance, the optics of a given experiment can be asymmetric; the detector gain may be unknown and time-dependent; the data may exhibit resonant frequencies due to electronics or cooling non-idealities; and the observations are invariably noisy. All these non-idealities must be properly understood before one can attempt to extract cosmology from the observations.

In this paper we consider one specific component within this global calibration problem, namely how to

estimate the statistical properties of the instrumental noise in light of real-world complications. This problem has of course already been addressed repeatedly in the literature (e.g., Prunet et al. 2001; Hinshaw et al. 2003), and our method is in principle similar to that of Ferreira & Jaffe (2000), taking a Bayesian approach to the problem. The main difference is that we formulate the algorithm explicitly in terms of a Gibbs sampler including both the time stream and the noise parameters as unknown variables, and this has several distinct advantages. First, it allows us to obtain proper uncertainties on all derived quantities. Second, gap filling is directly supported through built-in proper constrained realizations. This can for instance be used to account for instrumental glitches in the time stream, or to exclude point sources and other bright sources from the analysis. Third, it is straightforward to add support for additional nuisance parameters, due to the conditional nature of the Gibbs sampler. In this paper we implement template marginalization, which may for instance be useful for removing cosmic ray glitches in the Planck HFI data (Planck 2011b) or ground pickup for ground based experiments (e.g., QUIET 2011). We also outline the formalism for marginalization over CMB fluctuations, which may be relevant for experiments with high signal-to-noise ratio.

The method presented here is mathematically identical to the CMB Gibbs sampling framework developed by Jewell et al. (2004); Wandelt et al. (2004); Eriksen et al. (2004, 2008), and the main difference is simply that the object under consideration is a one-dimensional time stream instead of a two-dimensional field on the sphere. This makes the implementation considerably simpler, and the run times correspondingly faster. As a demonstration of the practicality of the method, we apply it to simulated data with properties typical for current ground-based experiments, and demonstrate explicitly that the computational costs of the method are tractable. The experiment of choice will be QUIET

i.k.wehus@fys.uio.no
 sigurdkn@astro.uio.no
 h.k.k.eriksen@astro.uio.no

¹ Theoretical Physics, Imperial College London, London SW7 2AZ, UK

² Department of Physics, University of Oslo, P.O. Box 1048 Blindern, N-0316 Oslo, Norway

³ Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029 Blindern, N-0315 Oslo, Norway

⁴ Centre of Mathematics for Applications, University of Oslo, P.O. Box 1053 Blindern, N-0316 Oslo, Norway

(2011), for which this method was initially developed.

2. DATA MODEL

The first step of any Bayesian analysis is to write down an explicit parametric model for the observations in question. In this paper, we start with the assumption that the output, \mathbf{d} , from a given detector can be written in terms of the following sum,

$$\mathbf{d} = \mathbf{n} + \mathbf{P}\mathbf{s} + \mathbf{T}\mathbf{a} + \mathbf{m}. \quad (1)$$

Here each term indicates a vector of n values sampled regularly in time in steps of Δt , that is, $\mathbf{d} = \{d_i\}$ with $i = 1, \dots, N$.

The first term on the right-hand side, \mathbf{n} , indicates the instrumental noise, which is our primary target in this paper. All the other components are only nuisance variables that we want to marginalize over.

We assume that the noise is Gaussian distributed and stationary over the full time range considered. In practice this means that the full data set of a given experiment should be segmented into parts which are individually piecewise stationary. For QUIET this corresponds to division into so-called “constant elevation scans” (QUIET 2011), while for Planck it corresponds to division into so-called “rings”, which are one-hour observation periods with a fixed satellite spin axis (Planck 2011a). Because the noise is assumed stationary, the time-domain noise covariance matrix, \mathbf{N} , depends only on the time lag between two observations, $N_{tt'} = N(t - t')$: It is a Toeplitz matrix, and may therefore be well approximated in Fourier domain with a simple diagonal matrix, $N_{\nu\nu'} = N_\nu \delta_{\nu\nu'}$. Here N_ν is the Fourier-domain noise power spectrum, which is given by the Fourier transform of $N(t - t')$.

Our main task is to estimate N_ν , and we do so in terms of a parametrized function. For many experiments this function is well approximated by a so-called $1/f$ profile,

$$N_\nu = \sigma_0^2 \left[1 + \left(\frac{\nu}{f_{\text{knee}}} \right)^\alpha \right] \quad (2)$$

which describes a sum of a correlated and an uncorrelated noise component in terms of three free parameters. The white-noise RMS level, σ_0 , defines the overall amplitude of the noise; the knee frequency, f_{knee} , indicates where the correlated and the uncorrelated components are equally strong, and α is the spectral index of the correlated component. Collectively, we denote $\{\sigma_0, \alpha, f_{\text{knee}}\}$ by θ . Of course, other parametrizations may easily be implemented if necessary.

The second term on the right-hand side, $\mathbf{P}\mathbf{s}$, indicates the contribution from the CMB sky, with \mathbf{P} being a pointing matrix, typically equal to zero everywhere except at P_{ip} if the detector points towards pixel p at time i , and s_p is the true (beam convolved) CMB signal. We make the usual assumption that \mathbf{s} is isotropic and Gaussian distributed with a given angular power spectrum, C_ℓ . In this paper, we will simply outline the formalism for how to deal with this term, and leave the implementation for a future paper dedicated to Planck analysis; as mentioned in the introduction, this machinery was initially developed QUIET, which is strongly noise dominated for a single data segment, and the CMB component is therefore not important, as will be explicitly

demonstrated in this paper.

The third term is a sum over n_{temp} time-domain templates. These can be used to model several different types of nuisance components. Three examples are diffuse foregrounds and cosmic ray glitches for Planck, and ground pick-up for QUIET. In either case, we assume in this paper that the template itself is perfectly known, and the only free parameter is an overall unknown multiplicative amplitude a . This is a vector of length n_{temp} , and \mathbf{T} is the two-dimensional $n \times n_{\text{temp}}$ matrix listing all templates column-wise.

Finally, the fourth term on the right-hand side of Equation 1 denotes a time-domain mask, \mathbf{m} . This is implemented by a “Gaussian” component having zero variance for samples that are not masked, and infinite variance for samples that are masked. In order to make analytic calculations more transparent, we write the corresponding covariance matrix as a diagonal matrix with elements

$$M_{ii} = \begin{cases} a, & i \text{ not masked} \\ \epsilon, & i \text{ masked} \end{cases}, \quad (3)$$

where $a \rightarrow \infty$ and $\epsilon \rightarrow 0$. A typical application of this component is to remove periods of instrumental glitching, or to discard particularly bright observations when the telescope points towards bright astrophysical sources, such as point sources or the Galactic plane.

3. GIBBS SAMPLING AND THE POSTERIOR

Our primary goal is now to map out $P(\theta|\mathbf{d})$, the noise spectrum posterior distribution marginalized over all nuisance components. By Bayes’ theorem this distribution reads

$$P(\theta|\mathbf{d}) = \frac{P(\mathbf{d}|\theta)P(\theta)}{P(\mathbf{d})} \propto \mathcal{L}(\theta)P(\theta), \quad (4)$$

where $\mathcal{L}(\theta) = P(\mathbf{d}|\theta)$ is the likelihood, $P(\theta)$ is a prior on θ , and $P(\mathbf{d})$ is an irrelevant normalization constant.

In this paper we adopt for simplicity uniform priors on σ_0 , α and f_{knee} . For typical relevant time series which contain $\sim 10^5$ samples, these parameters are usually strongly data-driven, and the choice of priors is therefore irrelevant. However, if an informative prior (or the Jeffreys’ prior) is desired for a given application, it is straightforward to include this as indicated by Equation 4.

Since we assume that the noise is Gaussian distributed with covariance $\mathbf{N}(\theta)$, the likelihood is given by

$$\mathcal{L}(\theta) \propto \frac{e^{-\frac{1}{2}\mathbf{n}^T\mathbf{N}^{-1}(\theta)\mathbf{n}}}{\sqrt{|\mathbf{N}(\theta)|}}, \quad (5)$$

where $\mathbf{n} = \mathbf{d} - \mathbf{P}\mathbf{s} - \mathbf{T}\mathbf{a} - \mathbf{m}$ is the noise component of the data stream. The goal is to compute this distribution, marginalized over \mathbf{s} and \mathbf{a} , while at the same time taking into account possible gaps in the data.

The latter point touches on an important computational issue. If there are no gaps in the data, then \mathbf{N} is a Toeplitz matrix, and multiplication with \mathbf{N} is performed most efficiently in Fourier space. However, the same does not hold if there are gaps in \mathbf{d} , since the symmetry of \mathbf{N} is broken. The well-known solution to this problem is to fill the gap with a constrained noise realization with the appropriate spectrum (e.g., Hoffman & Ribak 1991). In

our formulation, this is equivalent to estimating \mathbf{n} jointly with θ .

More generally, we want to estimate the joint density $P(\mathbf{n}, \mathbf{m}, \mathbf{s}, \mathbf{a}, \theta | \mathbf{d})$, from which any desired marginal may be obtained. At first sight, this appears like a formidable computational problem, involving more than 10^5 free parameters. However, this is also a problem that may be tackled by means of the statistical technique called Gibbs sampling, which has already been described in detail for computing the CMB angular power spectrum with contaminated data by Jewell et al. (2004); Wandelt et al. (2004); Eriksen et al. (2004, 2008).

According to the theory of Gibbs sampling, samples from a joint distribution may be obtained by iteratively sampling from each corresponding conditional distribution. For our case, this leads to the following sampling scheme,

$$\mathbf{m}, \mathbf{n} \leftarrow P(\mathbf{m}, \mathbf{n} | \mathbf{s}, \mathbf{a}, \theta, \mathbf{d}) \quad (6)$$

$$\mathbf{s}, \mathbf{n} \leftarrow P(\mathbf{s}, \mathbf{n} | \mathbf{a}, \theta, \mathbf{m}, \mathbf{d}) \quad (7)$$

$$\mathbf{a}, \mathbf{n} \leftarrow P(\mathbf{a}, \mathbf{n} | \theta, \mathbf{m}, \mathbf{s}, \mathbf{d}) \quad (8)$$

$$\theta \leftarrow P(\theta | \mathbf{n}, \mathbf{m}, \mathbf{s}, \mathbf{a}, \mathbf{d}) \quad (9)$$

The symbol \leftarrow indicates sampling from the distribution on the right-hand side. With this algorithm, $(\mathbf{n}, \mathbf{m}, \mathbf{s}, \mathbf{a}, \theta)^i$ will be drawn from the correct joint distribution.

Note that each of the sampling steps that involve time-domain vectors are joint steps including the noise component itself. This approach is highly computationally advantageous as it allows for fast multiplication with \mathbf{N} in Fourier domain; conditional algorithms for sampling each component separately would require slow convolutions in time domain. Of course, it is fully acceptable within the Gibbs sampling machinery to sample some components more often than others.

Note also that if we are only interested in the joint maximum-posterior parameters, we can replace the relevant steps in the above algorithm by a maximization operation, such that we maximize the conditional instead of sampling from it. The algorithm then reduces to a typical iterative approach, but formulated in a convenient and unified statistical language. The advantage of this approach is computational speed, while the disadvantage is the loss of information about uncertainties. Both versions of the algorithm will be implemented and demonstrated in the following.

4. SAMPLING ALGORITHMS

Equations 6–9 defines the high-level algorithm in terms of conditional sampling steps. To complete the method, we have to establish efficient sampling algorithms for each conditional distribution.

4.1. Noise estimation with ideal data

Perhaps the most fundamental conditional distribution in the sampling scheme outlined above is $P(\theta | \mathbf{n}, \mathbf{m}, \mathbf{s}, \mathbf{a}, \mathbf{d})$. This describes the distribution of the noise parameters given perfect knowledge about all components of the data. To obtain an explicit expression for this distribution, we first note that $P(\theta | \mathbf{n}, \mathbf{m}, \mathbf{s}, \mathbf{a}, \mathbf{d}) = P(\theta | \mathbf{n})$; if we know the true noise component, \mathbf{n} , no further information about either the CMB signal, the tem-




FIG. 1.— Correlation function for α of the Metropolis sampler employed to sample from $P(\theta | \mathbf{n})$ with a curvature matrix based proposal density. Similar plots for σ_0 and f_{knee} look visually the same. The correlations fall below 10% at a lag of ~ 20 samples, and we adopt a thinning factor of ~ 20 samples to suppress correlations, given that the computational cost of this sampling step is lower than the constrained realization sampler.

plate amplitudes, or, indeed, the actual data is needed in order to estimate the noise parameters.

The expression for the conditional distribution $P(\theta | \mathbf{n})$ is then formally the same as that for $P(\theta | \mathbf{d})$ given by Equations 4 and 5. Writing this out in Fourier space for the $1/f$ profile discussed in Section 2, one finds the following explicit distribution for $\theta = \{\sigma_0, \alpha, f_{\text{knee}}\}$,

$$-\ln P(\sigma_0, \alpha, f_{\text{knee}} | \mathbf{n}) = -\ln P(\sigma_0, \alpha, f_{\text{knee}}) + \sum_{\nu > 0} \left[\frac{p_\nu}{N_\nu} + \ln N_\nu \right]. \quad (10)$$

Here p_ν are the power spectrum components of the data \mathbf{n} , while $N_\nu = N_\nu(\sigma_0, \alpha, f_{\text{knee}})$ is the covariance matrix which in Fourier space is diagonal and given by Equation 2. The first term on the right-hand side is a user-defined prior.

To sample from this distribution, we use a standard Metropolis sampler with a Gaussian proposal density (e.g., Liu 2001). Each chain is initialized at the maximum-posterior point, which is found by a non-linear quasi-Newton search, and the covariance matrix of the Gaussian proposal density is taken to be the square root of the curvature matrix, evaluated at the maximum-posterior point. The elements of the inverse curvature matrix, $\mathcal{C}^{-1} = -\partial^2 \log P(\theta | \mathbf{n}) / \partial \theta_i \partial \theta_j$, read

$$\mathcal{C}_{\theta_i \theta_j}^{-1} = \sum_{\nu > 0} \left[\left(\frac{1}{N_\nu} - \frac{p_\nu}{N_\nu^2} \right) \frac{\partial^2 N_\nu}{\partial \theta_i \partial \theta_j} + \left(\frac{2p_\nu}{N_\nu^3} - \frac{1}{N_\nu^2} \right) \frac{\partial N_\nu}{\partial \theta_i} \frac{\partial N_\nu}{\partial \theta_j} \right]. \quad (11)$$

We have found that this proposal density leads to a Markov chain correlation length of about 20 samples for typical parameter values, and we therefore thin our chains by this amount. In addition, we remove a few post-thinned samples at the beginning of the chain to remove potential burn-in, although we have never seen evidence of any such effects. Finally, if we only want to find the maximum-posterior point, and not run a full-blown




FIG. 2.— Constrained realizations through a gap in the data. The extent of the gap is indicated by the two vertical dashed lines; the solid black curve shows the input data, which is dominated by a point source (thin spikes) in the masked out region. Five different constrained realizations are shown by colored lines. These depend only on the data outside the mask, and are therefore not affected by the point source.

Gibbs chain, as discussed in Section 3, we terminate the process directly after the initial quasi-Newton search.

4.2. Gap filling by constrained realizations

Next, we need to establish a sampling algorithm for $P(\mathbf{n}, \mathbf{m} | \mathbf{s}, \mathbf{a}, \theta, \mathbf{d}) = P(\mathbf{n}, \mathbf{m} | \mathbf{r}, \theta)$, where the residual $\mathbf{r} = \mathbf{d} - \mathbf{P}\mathbf{s} - \mathbf{T}\mathbf{a} = \mathbf{n} + \mathbf{m}$. Note that \mathbf{r} and \mathbf{n} differ only by \mathbf{m} , which represents contaminated segments of data that are masked out. The problem is therefore reduced to sampling the components of a sum of Gaussians given the sum itself. This may be achieved efficiently by solving the equation

$$[\mathbf{N}^{-1} + \mathbf{M}^{-1}] \mathbf{n} = \mathbf{M}^{-1} \mathbf{r} + \mathbf{N}^{-\frac{1}{2}} \mathbf{v} + \mathbf{M}^{-\frac{1}{2}} \mathbf{w} \quad (12)$$

for \mathbf{n} , where \mathbf{v} and \mathbf{w} are vectors of standard $N(0, 1)$ Gaussian random variates (e.g., Jewell et al. 2004; Wandelt et al. 2004; Eriksen et al. 2004).

The matrices involved here are typically of the order $10^5 \times 10^5$, so the equation cannot be solved by brute force. But since \mathbf{M} is diagonal in time domain and \mathbf{N} is diagonal in frequency domain (due to its Toeplitz nature), all multiplications are cheap in either time or Fourier domain, and the equation can be efficiently solved with the Conjugate Gradients method, properly changing basis as needed.

However, there is one practical complication involved in Equation 12: Since we intend to let $\epsilon \rightarrow 0$ and $a \rightarrow \infty$, the matrix $\mathbf{N}^{-1} + \mathbf{M}^{-1}$ becomes infinitely poorly conditioned. We can solve this problem, as well as significantly simplifying the equation, by splitting it into one equation for the masked region, and one for the unmasked region. Introducing the notation \mathbf{x}_1 and \mathbf{x}_2 for the unmasked and masked subsets of a vector \mathbf{x} respectively, we find

$$(\mathbf{N}^{-1} \mathbf{n})_1 + \epsilon^{-1} \mathbf{n}_1 = \epsilon^{-1} \mathbf{r}_1 + (\mathbf{N}^{-\frac{1}{2}} \mathbf{v})_1 + \epsilon^{-\frac{1}{2}} \mathbf{w}_1 \quad (13)$$

$$(\mathbf{N}^{-1} \mathbf{n})_2 + a^{-1} \mathbf{n}_2 = a^{-1} \mathbf{r}_2 + (\mathbf{N}^{-\frac{1}{2}} \mathbf{v})_2 + a^{-\frac{1}{2}} \mathbf{w}_2 \quad (14)$$

which in the limit $\epsilon \rightarrow 0$, $a \rightarrow \infty$ simplifies to

$$\mathbf{n}_1 = \mathbf{r}_1 \quad (15)$$

$$(\mathbf{N}^{-1} \mathbf{n})_2 = (\mathbf{N}^{-\frac{1}{2}} \mathbf{v})_2 \quad (16)$$

Note that equations 15–16 form an asymmetric equation system⁵. The Conjugate Gradients method is therefore not directly applicable, and the more general Biconjugate Gradients method must be used instead.

We found a simple diagonal preconditioner with value $\text{Var}(\mathbf{n})$ inside the mask and value 1 outside it to be sufficient.

4.3. Marginalization over fixed templates

The sampling algorithms described in Sections 4.1 and 4.2 defines the core noise Gibbs sampler, and together form a well-defined and complete noise estimation algorithm for low signal-to-noise data with gaps. However, one of the main advantages of the Gibbs sampling algorithm compared to other alternatives is its natural support for marginalization over nuisance parameters. In this section we describe the sampling algorithm for marginalization over fixed time-domain templates, describing for instance ground pickup, diffuse foregrounds or cosmic ray glitches.

First, we note that while the original data stream, \mathbf{d} , may contain gaps, and the Toeplitz nature of the noise covariance matrix is in that case broken, the constrained realization produced in Section 4.2 restores the Toeplitz symmetry. It is therefore computationally advantageous to use $\mathbf{q} = \mathbf{d} - \mathbf{P}\mathbf{s} - \mathbf{m} = \mathbf{n} + \mathbf{T}\mathbf{a}$ as the data in this step⁶, so that $P(\mathbf{a}, \mathbf{n} | \theta, \mathbf{m}, \mathbf{s}, \mathbf{d}) = P(\mathbf{a}, \mathbf{n} | \mathbf{q}, \theta)$.

Starting with Equation 5, solving for \mathbf{a} and completing the square in the exponential, the appropriate conditional distribution for \mathbf{a} is found to be the well-known distribution

$$P(\mathbf{a} | \mathbf{q}, \theta) \propto e^{-\frac{1}{2}(\mathbf{a} - \hat{\mathbf{a}})^t (\mathbf{T}^t \mathbf{N}^{-1} \mathbf{T})(\mathbf{a} - \hat{\mathbf{a}})}, \quad (17)$$

where $\hat{\mathbf{a}} = (\mathbf{T}^t \mathbf{N}^{-1} \mathbf{T})^{-1} \mathbf{T}^t \mathbf{N}^{-1} \mathbf{q}$; that is, $P(\mathbf{a} | \mathbf{q}, \theta)$ is a Gaussian distribution with mean $\hat{\mathbf{a}}$ and covariance $\mathbf{C}_a = (\mathbf{T}^t \mathbf{N}^{-1} \mathbf{T})^{-1}$. The same result has been derived for numerous other applications, one of which was described by Eriksen et al. (2004), outlining template amplitude sampling with CMB sky map data.

Sampling from this distribution is once again straightforward. If one only wish to marginalize over a small number of templates, the easiest solution is simply to compute both $\hat{\mathbf{a}}$ and \mathbf{C}_a by brute-force, and let $\mathbf{a}^{i+1} = \hat{\mathbf{a}} + \mathbf{L}_a \eta$, where \mathbf{L}_a is the Cholesky factor of $\mathbf{C}_a = \mathbf{L}_a \mathbf{L}_a^t$, and η is a vector of uncorrelated standard Gaussian $N(0, 1)$ variates. On the other hand, if there are more than, say, 1000 templates involved, it may be faster to

⁵ Equations 15–16 can be rewritten as

$$\begin{bmatrix} 1 & 0 \\ N_{21}^{-1} & N_{22}^{-1} \end{bmatrix} = \begin{bmatrix} \mathbf{r}_1 \\ (\mathbf{N}^{-\frac{1}{2}} \mathbf{v})_2 \end{bmatrix}$$

This is asymmetric because we multiplied the unmasked part of the equation by ϵ in order to get a finite result.

⁶ Of course, one could write down a sampling algorithm for \mathbf{a} that only uses the non-masked data directly, but this would require heavy time-domain convolutions, and not take advantage of the symmetries inherent in the noise covariance matrix.




FIG. 3.— Recovered noise parameters from 10 000 noise-only simulations with gaps. The dashed horizontal lines indicate the true input value.

solve the following equation by Conjugate Gradients,

$$(\mathbf{T}^t \mathbf{N}^{-1} \mathbf{T}) \mathbf{a} = \mathbf{T}^t \mathbf{N}^{-1} \mathbf{q} + \mathbf{T}^t \mathbf{N}^{-\frac{1}{2}} \omega, \quad (18)$$

where ω is a full time-stream of $N(0, 1)$ random variates. In this paper, which has the QUIET experiment as its main application, we only use the former algorithm; for Planck the latter algorithm may be useful in order to account for frequent and partially overlapping cosmic ray glitches efficiently.

4.4. CMB signal marginalization

Most CMB experiments are strongly noise-dominated within relatively short time periods, and need to integrate over the sky for a long time in order to produce high-sensitivity maps. For instance, for the first-season QUIET experiment the mean polarized sensitivity of a detector was $280 \mu\text{K}\sqrt{\text{s}}$ (QUIET 2011), while the polarized CMB sky has an RMS of $\sim 1 \mu\text{K}$ on the relevant angular scales. In such cases, it is an excellent approximation simply to ignore the CMB contribution when estimating the noise parameters. However, this approximation does not hold for all experiments, and one particularly important counterexample is Planck.

In order to marginalize over the CMB signal we need to be able to sample from $P(\mathbf{s}, \mathbf{n} | \mathbf{a}, \theta, \mathbf{m}, \mathbf{d}) = P(\mathbf{s}, \mathbf{n} | \mathbf{u}, \theta)$, where the residual is $\mathbf{u} = \mathbf{d} - \mathbf{m} - \mathbf{F}\mathbf{a} = \mathbf{P}\mathbf{s} + \mathbf{n}$. This is again a case that involves sampling terms of a Gaussian sum given the sum itself. The only difference from Section 4.2 is that a projection operator is involved in this case,

$$(\mathbf{S}^{-1} + \mathbf{P}^T \mathbf{N}^{-1} \mathbf{P}) \mathbf{s} = \mathbf{P}^T \mathbf{N}^{-1} \mathbf{u} + \mathbf{S}^{-\frac{1}{2}} \mathbf{v} + \mathbf{P}^T \mathbf{N}^{-\frac{1}{2}} \mathbf{w} \quad (19)$$

Here \mathbf{v} and \mathbf{w} are vectors of standard normal samples in pixel and time domain respectively. Note that \mathbf{M} is not involved this time, so the expression can be used as it is. As before, multiplications involving \mathbf{N} are efficient in Fourier space due to the Toeplitz nature of \mathbf{N} , but the CMB signal covariance, \mathbf{S} , will in general be dense, depending on the true CMB power spectrum. This makes this method prohibitively expensive for general asymmetric scanning strategies. There are, however, circumstances for which also this multiplication becomes efficient. One obvious case is that of full sky coverage,

where a change to spherical harmonic basis makes \mathbf{S} diagonal.

More interestingly, \mathbf{S} also becomes diagonal when expressed in Fourier basis on a circle on the sky. To see this, consider a circle with radius θ on the sphere, parametrized by the angle ϕ . The covariance between the points $p_1 = (\theta, \phi_1)$ and $p_2 = (\theta, \phi_2)$ with angular distance r is given by the two-point correlation function:

$$S(p_1, p_2) = C(r) = \frac{1}{4\pi} \sum_{l=0}^{\infty} (2l+1) C_l P_l(\cos(r)) \quad (20)$$

Since $r(p_1, p_2)$ for the case of a circle only depends on $\Delta\phi$, and is independent of ϕ itself, $S(\phi_1, \phi_2)$ is a Toeplitz matrix, and is therefore diagonal in Fourier space. This is highly relevant for Planck, since the Planck scanning strategy (Planck 2011a) naturally divides into scans of circles. Therefore, in this case sampling $\mathbf{s} \leftarrow P(\mathbf{s} | \mathbf{u})$ can be done at very low extra cost.

5. APPLICATION TO SIMULATED DATA

In this section we demonstrate the noise Gibbs sampler as described in Section 4 on a particular type of QUIET simulations. QUIET is a radiometer-based CMB B-mode polarization experiment located in the Atacama desert (QUIET 2011), which took observations from August 2008 to December 2010. The first results were based on only nine months worth of data, and yet already provided the second most stringent upper limit on the tensor-to-scalar ratio, r , to date based on CMB polarization measurements.

In its normal mode of operation, QUIET observed four separate CMB fields on the sky which were each chosen because of their low foreground levels. In this mode, the experiment is totally noise dominated on time scales of an hour or less, with a mean polarized sensitivity per detector diode of $280 \mu\text{K}\sqrt{\text{s}}$ (QUIET 2011).

However, QUIET also observed two Galactic patches, one of which was the Galactic center, as well as several bright calibration objects, such as the Moon, Tau Alpha and RCW38. These sources are bright enough to be seen visually in each detector time stream, and they can therefore bias any noise estimates unless properly accounted for. Such objects also complicate automated data selection processes, since it is difficult to distinguish between

TABLE 1
VALIDATION BY SIMULATIONS

Simulation	σ_0 (10^{-5}V)	α	f_{knee} (10^{-1}Hz)
Posterior maximization (absolute parameter values)			
Gaps only	1.000 ± 0.003	-1.80 ± 0.06	1.00 ± 0.05
Gaps + uncorrected CMB	1.000 ± 0.003	-1.80 ± 0.06	1.00 ± 0.05
Gaps + uncorrected ground pickup	1.006 ± 0.003	-1.52 ± 0.05	1.51 ± 0.08
Gaps + corrected ground pickup	1.000 ± 0.003	-1.80 ± 0.06	1.00 ± 0.05
Gibbs sampling (normalized parameter values)			
Gaps only	-0.02 ± 1.00	0.01 ± 1.01	0.03 ± 1.01
Gaps + uncorrected CMB	-0.02 ± 1.00	0.00 ± 1.01	0.02 ± 1.01
Gaps + uncorrected ground pickup	1.78 ± 1.03	6.04 ± 1.36	6.63 ± 0.77
Gaps + corrected ground pickup	-0.02 ± 1.00	0.00 ± 1.00	0.03 ± 1.00

NOTE. — Summary of recovered noise parameters from various simulated ensembles. Each column indicates the mean and standard deviation of the resulting parameter distribution. The top section shows results obtained when simply maximizing the posterior, while the bottom section shows the results for a full Gibbs sampling analysis. Each run in the top section consists of 10000 simulations, while each run in the bottom section consists of 5000 simulations. All runs have been started with the same seed, to ensure directly comparable results.

an astrophysical object and an instrumental glitch.

In this section we show how the algorithm developed in Section 4 may be applied to such situations. Specifically, we consider a observing session lasting for about 40 minutes of a field including a bright source with known location, and assume that the data may be modeled as $\mathbf{d} = \mathbf{n} + \mathbf{T}\mathbf{a} + \mathbf{m}$. Here \mathbf{T} is a single template describing possible sidelobe pick-up from the ground, constructed from the full observing season as described by QUIET (2011), and \mathbf{m} is a time-domain mask that removes any samples that happen to fall closer than 1° from the source center. The total number of samples in the data stream is 60 949, and the total number of masked samples is 651.

The simulations used in this section are constructed as follows. We set up an ensemble of 10^4 time streams containing correlated Gaussian random noise with $\sigma_0 = 10^{-5}\text{V}$, $\alpha = -1.8$ and $f_k = 0.1\text{Hz}$; the white noise and spectral index are representative for a QUIET detector, while the knee frequency is grossly exaggerated to push the algorithm into a difficult region of parameter space, as well as to more clearly visualize the outputs of the algorithm. A far more reasonable value for QUIET is $f_k = 10\text{mHz}$, and we have of course verified that the algorithm also works for such cases. Further, it reaches convergence faster in that case than with the extreme value of f_k used in the present simulations.

5.1. Visual inspection of constrained realizations

Before considering the statistical properties of the resulting posterior distributions, it is useful to look visually at a few constrained realizations in order to build up intuition about the algorithm. In order to highlight the behavior of the constrained realizations, we make two adjustments to the above simulation procedure for this case alone: First, we replace the tuned mask with a wide 6000-sample mask, covering the entire time range in which the source is visible, and second, we make the correlated noise component stronger by setting $\sigma_0 = 10^{-6}\text{V}$, $f_{\text{knee}} = 1\text{Hz}$ and $\alpha = -2.3$.

The results are shown in Figure 2. The raw data are shown in the solid black line, and the vertical dashed

lines indicate the extent of the gap. The colored curves within the gap shows 5 difference constrained realizations; note that together with the black solid curve outside the mask, any of these form a valid noise realization with the appropriate noise power spectrum as defined by σ_0 , α and f_{knee} . They are each a valid sample drawn from $P(\mathbf{n}, \mathbf{m}|\mathbf{d})$. However, if one had tried to estimate the noise spectrum also using the data inside the gap, the source signal (seen as sharp spikes in Figure 2) would bias the resulting noise parameters.

In this paper, we consider the constrained realizations primarily to be a useful tool that allows for fast noise covariance matrix multiplications in Fourier space. However, these constrained realizations can of course also be useful in their own right, for instance for deglitching a time stream before map making.

5.2. Validation and statistical characterization

We now seek to statistically validate our algorithms and codes. Both the posterior maximization and the Gibbs sampling algorithms are considered. The number of simulations are 10 000 for posterior maximization and 5000 for Gibbs sampling, with properties as described above. In each case, we consider four different models. First, we analyze simulations including only noise and gaps. Second, we add a CMB signal to each realization, but do not attempt to correct for it. Third, we add a strong ground template to each realization, and do also not attempt to correct for it. Fourth, we analyze the same ground-contaminated simulations as above, but this time do marginalize over an appropriate template. The same random seeds were used in each of the four simulation and analysis results, in order to allow for direct comparison of results between runs.

The results from this exercise are summarized in Table 1, and histograms for the first case are shown in Figure 3. First, the the upper section in Table 1 lists the mean and standard deviation of the recovered parameters for the posterior maximization algorithm. Recall that the input parameters were $(\sigma_0, \alpha, f_{\text{knee}}) = (10^{-5}\text{V}, -1.8, 0.1\text{Hz})$, and these are recovered perfectly




FIG. 4.— *Left*: CPU time as a function of total number of samples in the time stream, keeping the number of masked samples constant. *Middle*: CPU time as a function of number of masked samples. The size of each gap is kept constant at 100 samples, and only the number of gaps vary. *Right*: CPU time as a function of number of masked samples, but in this case there is only one gap of varying width.

in all cases, except for the one involving an uncorrected ground template, as expected.

Second, in the bottom half we show the results from the Gibbs sampling analyses, but this time in terms of normalized parameters on the form $r = (\theta_{\text{est}} - \theta_{\text{in}})/\sigma_{\text{est}}$, where θ_{est} and σ_{est} are the mean and standard deviation of the Gibbs chain for a given parameter (removing the first 10% of the samples for burn-in), and σ_{in} is the true input value. If the Gibbs chain is both unbiased and has the correct dispersion, r should be Gaussian distributed with zero mean and unit variance. As seen in Table 1, this is indeed the case.

We also note that adding a CMB component to these simulations do not bias the noise estimates, simply because the CMB is too weak to be detected on the time scales considered here. This confirms the assumption made by the QUIET team when estimating their noise properties: The QUIET observations are sufficiently noise dominated on a one-hour time scale that the CMB can be safely neglected for noise estimation purposes.

5.3. Resource requirements

To be practical, it is not sufficient that a method is robust and accurate, but it must also be computationally efficient. For the present algorithm the two most important parameters for computational speed are 1) the total number of samples in the time stream, n , and 2) the number of masked samples, m , while also the relative position of the masked samples play an important role.

In Figure 4 we show the scaling of each of the two algorithms (posterior maximization and Gibbs sampling) as a function of both n (left panel) and m (right panel). In the left plot, m was fixed at 1000, divided into ten gaps of 100 samples each, and only the total length of the

data stream was varied. In this case, we should expect the scaling of the overall algorithm to be dominated by Fourier transforms, suggesting an overall behavior given by $\mathcal{O}(n \log n)$. As seen in Figure 4, this approximation holds to a very high degree, both for posterior maximization and Gibbs sampling. Further, we see that the CPU time required to analyze a single 100 000 sample data set with 1000 samples removed is 3 seconds for posterior maximization and 21 seconds for Gibbs sampling.

In the middle panel, we fix n at 100 000, and increase m by varying the number of gaps, each extending 100 samples. Perhaps somewhat surprisingly, we see that the computing time in this case is nearly independent of m . The reason for this is simply that the number of conjugate gradient iterations required for the gap filling procedure is largely determined by condition number (ie., the ratio between the highest and smallest eigenvalue) of the covariance matrix of a single gap. Having more gaps separated by more than one time-domain correlation length effectively corresponds to performing multiple matrix inversions in parallel, and the net cost therefore do not increase significantly.

In the third panel, we increase m by making one gap larger, as opposed to adding many small gaps. In this case the CPU time does increase dramatically, because it becomes increasingly hard for the algorithm to fill the missing pieces of the data stream. In this case, the noise covariance matrix condition becomes larger and larger.

6. SUMMARY

We have described and implemented a Bayesian framework for estimating the time-domain noise power spectrum for non-ideal CMB experiments. This framework is conceptually identical to a previously described method

for estimating the angular CMB power spectrum from CMB sky maps (Jewell et al. 2004; Wandelt et al. 2004; Eriksen et al. 2004), and relies heavily on the Gibbs sampling algorithm. The single most important advantage of this method over existing competitors in the literature derives from the conditional nature of the Gibbs sampler: Additional parameters may be introduced *conditionally* into the algorithm. This allows for seamless marginalization over nuisance parameters, which otherwise may be difficult to integrate. A second important advantage of the method is the fact that it provides proper uncertainties on all estimated quantities, which at least in principle later may be propagated into final estimates of the uncertainties of the CMB sky map and angular power spectra.

In this paper we implemented support for two general features that are useful for analysis of realistic data, namely constrained realizations and template sampling. The former is useful whenever there are gaps in the data, for instance due to an instrumental glitch, or there are strong localized sources in the sky that may bias the noise estimate: In these cases, the gaps are refilled with a constrained noise realization with the appropriate noise parameters, such that the full time stream represents a proper sample from a Gaussian distribution with a noise covariance matrix, \mathbf{N} . Since the time stream no longer contains gaps, the Toeplitz symmetry of the noise covari-

ance matrix is restored, and matrix multiplications may be performed quickly in Fourier space.

The second operation, template sampling, is also a powerful and versatile technique for mitigating systematic errors. In this paper we mostly focused on data from the QUIET experiment, for which ground pickup from sidelobes is one significant source of systematics (QUIET 2011). In a future publication we will apply the same method to simulations of the Planck experiment, for which cosmic ray glitches is an important source of systematic errors. As detailed by Planck (2011b), these cosmic rays may be modeled in terms of a limited set of time domain templates, and the algorithms presented in this paper should therefore prove useful for mitigating the effects of these glitches, as well as for propagating the corresponding uncertainties into the final noise spectrum parameters.

We thank the QUIET collaboration for stimulating discussions. The computations presented in this paper were carried out on Titan, a cluster owned and maintained by the University of Oslo and NOTUR. This project was supported by the ERC Starting Grant StG2010-257080. Some of the results in this paper have been derived using the HEALPix (Górski et al. 2005) software and analysis package.

REFERENCES

- Eriksen, H. K., et al. 2004, ApJS, 155, 227
 Eriksen, H. K., Jewell, J. B., Dickinson, C., Banday, A. J., Górski, K. M., & Lawrence, C. R. 2008a, ApJ, 676, 10
 Ferreira, P. G., & Jaffe, A. H. 2000, MNRAS, 312, 89
 Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., Bartelman, M. 2005, ApJ, 622, 759
 Hinshaw, G., et al. 2003, ApJS, 148, 63
 Hoffman, Y., & Ribak, E. 1991, ApJ, 380, L5
 Komatsu, E., et al. 2011, ApJS, 192, 18
 Jewell, J., Levin, S., & Anderson, C. H. 2004, ApJ, 609, 1
 Liu, J. S., Monte Carlo Strategies in Scientific Computing, Cambridge, USA: Springer, 2001,
 Planck Collaboration 2011, A&A, submitted, [1101.2022]
 Planck HFI Core Team 2011, A&A, submitted, [1101.2048]
 Prunet, S., et al. 2001, in MPA/ESO/MPA conf. Proc. "Mining the Sky", [astro-ph/0101073]
 QUIET Collaboration 2011, ApJ, in press, [1012.3191]
 Tauber, J. A., et al. 2010, A&A, 520, A1
 Wandelt, B. D., Larson, D. L., & Lakshminarayanan, A. 2004, Phys. Rev. D, 70, 083511